

A Complete framework for ambush avoidance in realistic environments

Emmanuel Boidot
eboidot3@gatech.edu

Aude Marzuoli
amarzuoli3@gatech.edu
School of Aerospace Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332-0250

Eric Feron
feron@gatech.edu

Abstract—Operating vehicles in adversarial environments between a recurring origin-destination pair requires new planning techniques. A Ruckle-inspired game is introduced. The goal of the first player is to minimize the expected casualties undergone by a convoy. The goal of the second player is to maximize this damage. The outcome of the game is obtained via a linear program that solves the corresponding minmax optimization problem over this outcome. Different environment models are defined in order to compute routing strategies over unstructured environments. To compare these methods for increasingly accurate representations of the environment, a grid-based model is chosen to represent the environment and the existence of a sufficient network size is highlighted. A global framework for the generation of realistic routing strategies between any two points is described. This framework requires a good assessment of the potential casualties at any location, therefore the most important parameters are identified. Finally the framework is tested on real world environments.

I. INTRODUCTION

We seek to develop routing strategies for dynamical systems operating in adversarial environments. For the problem at hand, an agent is trying to move a vehicle from a given origin a desired target set, while avoiding undesired set of states (ambushes), and another agent is rewarded if the system is caught in one of the ambushes. Examples of applications of this research could be found in protection of automated vehicles, such as delivery vehicles, or piracy prevention.

Recently, significant research has focused on dynamic pursuit-evasion games [9]. Unlike these efforts, our work is interested in computing mixed planning strategies. In the context of differential pursuit-evasion games, the solutions found are often found to be deterministic. Our problem returns a random distribution of solutions. This means that, for the same scenario, the returned strategy is unique but two vehicles following this strategy might use different paths. For clarity reasons, a probabilistic strategy (detailed later) will be described as a "route" and a deterministic realization of this strategy will be described as a "path".

The idea behind this research effort was first introduced by Ruckle [12], who extended Isaacs's [6] classical battleship versus bomber duel by interpreting a two-dimensional environment as a rectangular array of lattice points. Ruckle stated that such geometric games played on a finite lattice always exhibit a solution because of the Minimax theorem of Von Neumann [10]. He also identified necessary and sufficient conditions for

the mixed strategy of a player to be optimal. Ruckle's idea was later extended by Farmey and Feron [8] in order to formulate this game as a two player zero sum game applicable to a non-lattice network. While they advanced the idea of a variable game outcome at each node of the network, no analysis of this parameter was performed. Their formulation was later used in applications for piracy prevention [15][7] and money transit protection [14], where the multiple destination case is tackled.

The specific goal in our approach is to exit the purely theoretical world of game theory in order to perform simulations close to real-life applications. This is achieved through two distinct processes. First, a method to create a network on any environment on which the game is to be run is developed. Second, the outcome of an ambush is no longer constant over the environment. The outcome of an ambush now depends on its position and the factors influencing this outcome are identified. Therefore the optimal routing strategy can be obtained on any environment without manual input of a case specific game matrix. The authors believe that these two points are the main contribution of the present paper to the current state of the art in the field of ambush games.

The outline of this paper is as follows. In Section II the game and its mathematical formulation are further described. Section III details different network construction methods and ambush types considered for the adaptation of the game to unstructured environments, while Section IV discusses the results of the game for different methods and parameters. Section V describes the factors of risk for ground vehicles and illustrates the use of our work on several applications.

II. APPROACH

A. Game Description

The problem of interest is to plan a path for a convoy that needs to journey from an origin A to a destination B in an environment where hostile forces might try to ambush it. It is modelled as a two players non-cooperative zero-sum game, meaning that a player's sole goal is to optimize his personal gain (non cooperative) and the sum of the outcome over all players is zero. Player 1, BLUE, chooses a route from origin to destination. Player 2, RED, selects a number of locations at which to set up ambushes. The number of ambushes depends on RED's resources. In the remainder of

this paper RED only owns one resource, i.e. it can only set one ambush. This assumption can be made without consequences because we are interested in mixed strategies. It makes the sum of probability over all possible states for RED equal to one instead of a higher integer. If BLUE's path passes through an ambush site, then RED wins. If BLUE's path avoids all ambush sites, then BLUE wins. The outcome of an ambush corresponds to the casualties that the agent would experience in this specific area and it is dependent on the characteristics of the local environment. The outcome of a game is the sum of this local outcome over all the ambushes that BLUE's path has gone through. The outcome α of an ambush at point x satisfies $\alpha = f(x)$. It is assumed to be known by both players. Its approximation is discussed in Section V.

The environment is represented by a network $(\mathcal{N}, \mathcal{E})$ and a local outcome map. Each ambush area is associated with a local outcome α_i . The set α is the discretized local outcome map, which measures the possible losses for BLUE. In this paper only single-stage ambush games are considered, meaning that both players need to decide their strategy at the beginning of the game and cannot reconsider it during the game. Neither player has any information about the other player's current path. The environment and its discrete representation are the only information RED and BLUE share : the network, the local outcome map, BLUE's origin and destination points and the position of the ambush areas (but not the position of those chosen by RED) are known to both players.

A possible strategy for BLUE is represented by a probability vector \mathbf{p} , where p_{ij} is the probability that the convoy uses edge e_{ij} between nodes n_i and n_j . Similarly, a strategy for RED is represented by a probability vector \mathbf{q} that contains the probability q_k that RED sets an ambush at the ambush area a_k . The physical meaning of an ambush area corresponds to the area RED can act on if he sets his ambush inside a given area. Setting an ambush on this area means that BLUE will be ambushed if its path goes through nodes in this set. It is common for this type of game theoretic problems to assume that ambushes take place at nodes or on edges of the network, which is the approach adopted in previous work [3]. However, to expand the approach to continuous state space, it makes more practical sense for RED's strategies to relate to areas of the environment (continuous sets). This way, RED's strategy translates into where it would position its ambushes given a specified action range.

Given the network and the local outcome map, the goal of the game is to find the optimal strategy \mathbf{p}^* for BLUE, assuming that RED follows its optimal strategy \mathbf{q}^* . Consider the case of a recurring continuous transition from state A to state B, where there exist two different paths from A to B. A deterministic approach will return only one of the two paths, whereas the mixed strategy p will return a probability p_1 of using the first path and a probability p_2 of using the second path. Incorporating mixed strategy solutions means that, for a given environment, the opponent can at best (after a large number of runs following this strategy) figure out the strategy p but will never gather any certitude over the convoy's exact

path.

B. Mathematical Formulation

Let $(\mathcal{N}, \mathcal{E})$ be a network. The most simple example case where ambushes happen at nodes is considered first. Therefore ambush areas $\{a_i\}$ and nodes $\{n_i\}$ represent the same set in this subsection. Let \mathbf{p} (resp. \mathbf{q}) be the probability vector representative of BLUE's (resp. RED's) mixed strategy. Assume that the two players strategies are independent. At each node n_j , the probability that BLUE gets ambushed is equal to the probability that BLUE's path goes through n_j times the probability that RED sets an ambush at this node. The gain for RED at this node being α_i , the expected outcome of the game relative to this node is: $\sum_{i|(i,j) \in \mathcal{E}} p_{ij} q_j \alpha_j$. Therefore the strategic outcome of the game is :

$$\mathcal{V} = \sum_{j \in \mathcal{N}} \sum_{i|(i,j) \in \mathcal{E}} p_{ij} q_j \alpha_j = \mathbf{q}^t \mathbf{D} \mathbf{p}. \quad (1)$$

with $D_{jk} = \alpha_j$ if the k^{th} line of \mathbf{p} represents the probability that BLUE uses an edge e_{ij} directed towards n_j , and $D_{jk} = 0$ otherwise.

The objective of the approach is to find the strategy for BLUE that minimizes the maximal possible outcome for RED. Provided that $q_j \leq 1 \quad \forall j$, RED can always maximize \mathcal{V} by choosing the node n_j for which the probability of BLUE passing through that node weighted by the value α_j is maximal. Therefore BLUE's optimal solution is to minimize this product across all nodes :

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \left(\max_{j \in \mathcal{N}} \sum_{i|(i,j) \in \mathcal{E}} p_{ij} \alpha_j \right). \quad (2)$$

The other constraints of this problem enforce the flow conservation through the network. The probability that the convoy arrives at node n_j is equal to the probability that the convoy leaves the same node. Probabilities of the convoy being at origin and destination nodes are equal to 1. This problem is solved as a linear optimization problem by introducing a variable z constrained as follows.

$$z \geq \sum_{i|(i,j) \in \mathcal{E}} p_{ij} \alpha_j \quad \forall j \in \mathcal{N} \quad (3)$$

Rewriting Equations (2) and (3), the problem can be posed as a linear problem, with A and b representing the flow conservation constraints.

$$\begin{array}{ll} \text{minimize } & z \\ \text{subject to } & \mathbf{D} \mathbf{p} - \mathbf{1} z \leq \mathbf{0} \\ & \mathbf{A} \mathbf{p} = \mathbf{b} \\ & \mathbf{p} \geq \mathbf{0} \end{array} \quad (4)$$

III. FRAMEWORK FOR UNSTRUCTURED ENVIRONMENTS

Though the approach described in Section II has been partially tackled in [8, 2], the previous studies assumed prior existence of a network to optimize the routing on. Using this type of model on different environments requires investigating what a good representation of the environment is. This section

Method #	Sampling	Connectivity
1(rdm)	Random	Delaunay triangulation
2(uni8)	Uniform	8 connected grid
3(uniD)	Uniform	Delaunay triangulation

TABLE I
DIFFERENT NETWORK CONSTRUCTION METHODS.

presents three methods to build the network that supports the optimization. The outcome of the strategic optimization is compared for different representation of the environment and some key features are identified.

A. Network Construction

Assuming the existence of a network limits the number of possible routes for the vehicle. Therefore, it decreases the advantage of this approach which aims at increasing the variability of the convoy's trajectory, in particular for vehicles with off-road capabilities. Hence it is expected that expanding the structured network with a graph representation of the unstructured environment will improve the approach efficiency. This representation has to be adapted to the vehicle physical model and environment characteristics.

Several methods are tested in order to create a high fidelity directed network with respect to the environment while allowing reasonably short computation time. Differences were made on the sampling method and on the connectivity between nodes, as displayed in Table I. While randomly sampled nodes connected through a Delaunay triangulation result in a relatively small and computationally efficient representation of the environment, *rdm* might not be representative enough of the details of the environment. The most complete representation of the environment is obtained through *uni8*, but it requires 16 directed edges per node (8 neighboring nodes, two directions per edge). This method is computationally intensive, therefore it might be preferable to choose a less precise technique with yet enough precision on the environment description. These methods will be compared in Section IV.

B. Ambush types

A different model is now considered where the state space for RED is modified. Ambushes are paired with an area in the environment instead of a node on the network. Since the initial purpose of this framework was to create routes minimizing losses on existing networks, considering point-based ambushes was a starting point. Now, the case of fixed geographic ambush areas is considered. This model is more realistic in the sense that it relays the idea of a maximum reach for RED, which defines the size of the ambush areas. Another important reason to have ambush by area is that it allows to decouple the space dependency for BLUE strategy from RED's strategy in order to study the convergence of BLUE's strategy regarding the precision of the environment model when the set of ambushes for RED is fixed. This is a very interesting feature if one

wishes to use the discrete routing strategies as approximations of a continuous strategy. The results will give us insights on the possibility to construct such a continuous solution.

For this case, a transposition matrix \mathbf{S} is created. \mathbf{S} describes the belonging of a node to an ambush area: if node j belongs to ambush area i then the j^{th} column of \mathbf{S} will be zero except on its i^{th} line. This notation allows the reformulation of the linear problem to be very close to the initial formulation. BLUE's optimal solution is now: $\mathbf{p}^* = \arg \min_{\mathbf{p}} \max_{\mathbf{q}} \mathbf{q}^t \mathbf{S} \mathbf{D} \mathbf{p}$.

Looking back at the trivial example, two examples of ambush area sets are displayed in Figure 1. As reflected in this Figure, the reach of RED has a great influence on BLUE's optimal strategy. This validates our motivation to include this parameter in the environment model. Note that the strategy \mathbf{q}^* for RED is computed assuming it has perfect knowledge of BLUE's strategy. This means that RED knows the network and the probability that BLUE uses any edge of the network. Once \mathbf{p}^* is computed for BLUE, \mathbf{q}^* is computed as $\mathbf{q}^* = \arg \max_{\mathbf{q}} \mathbf{q}^t \mathbf{S} \mathbf{D} \mathbf{p}^*$. This is a very strong assumed advantage for RED regarding which this routing method is quite robust.

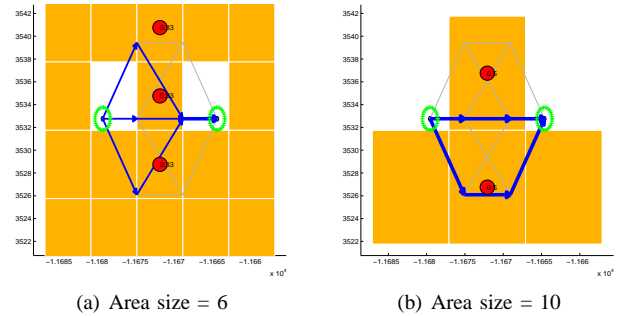


Fig. 1. Simple network example - Ambushes by area. The probability associated with each edge is displayed as the width of the edge. The orange rectangles represent the intensity of the local outcome on corresponding ambush areas. The origin and destination area are supposed to have a zero local outcome. Grey edges are unused or used with very low probability. Red circles represent the strategy chosen by RED, ie areas i where RED might set his ambush with probability $q_i > 0$. This example illustrates the importance of associating an area to an ambush. Increasing the reach of RED between (a) and (b) leads to a drastically different strategy: the reach of RED has very important consequences on BLUE's strategy.

C. Energy Optimization

A consequence of taking into account the reach of RED is the creation of cycles inside each ambush area. This means that the probabilistic routing could lead to a path where the vehicle stays inside a closed subspace for a very long time. This situation is more likely in ambush areas where the local outcome is close to zero. This illustrates one of the issues with this routing method, which is the absence of any kind of energy optimization.

In order to avoid cycles, an energy optimization factor λ is introduced that allows to tune the outcome of the optimization. Large values of λ will lead to a few different paths close to the shortest path while increasing the overall risk for the

system. Low values of λ will return routing strategies close to the safest ones. Note that, the set of optimal strategy being quite large, the routing strategies returned with low values of λ might also be risk optimal in several situations. The new objective function is $(1 - \lambda)z + \lambda \sum_{i,j|(i,j) \in \mathcal{E}} p_{ij} \|e_{ij}\|$.

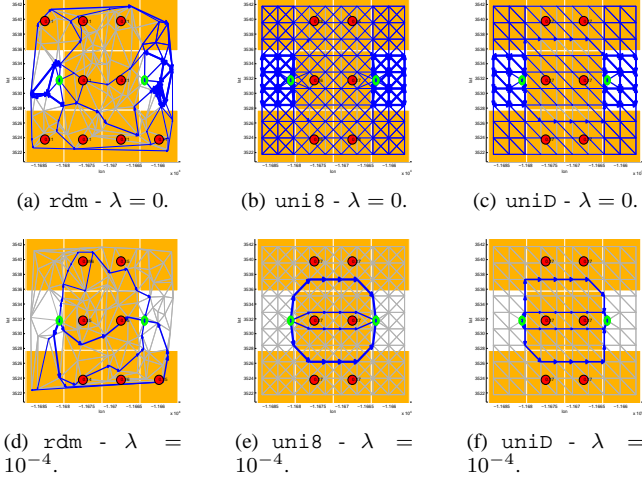


Fig. 2. Results of different network construction methods. In the left column, cycles can be identified in the departure and arrival areas whereas there are only outgoing edges in the right column examples. An interesting property of the route proposed with energy optimization is the 1-to-1 repartition of route segments to areas (as seen in Figures 2(e) and 2(f)) along the median between departure and arrival.

The optimal routes returned by our framework for different network construction method can be first analyzed through their graphical representation. Examples of routing strategies with low or no energy optimization for different network construction methods described in III-A are compared in Figure 2. It is important to note that the routes with and without energy optimization are very much alike. Therefore the cycles are suppressed without significant change to the routing strategy regarding the outcome of the game. The energy optimization also allows to avoid a side effect where flows agglomerate on the border of the environment.

Interestingly, it seems that the route includes a fixed number of distinct paths depending on the reach of RED. In Figure 2, this distance is such that the median between departure and arrival is divided into three segments. This leads to four distinct paths in Figures 2(f), 2(e). Yet, this is optimally equivalent to having only one path along the set of areas aligned between departure and arrival. In fact any route where paths do not intersect with a uniform flow distribution along the ambush areas would be optimally equivalent regarding the strategic outcome. Finally the network construction method seem to have an important impact on RED strategy. Comparing Figures 2(a), 2(d) and Figures 2(b), 2(e), 2(c), 2(f), RED has a strategy with a larger spread on grid-based networks. This could lead us to conclude that these networks allow for better results. This is however not the case because the strategy displayed in Figures 2(b), 2(e), 2(c), 2(f) is similar to a single ambush area strategy with probability 1.0 across any area on

the departure-arrival median.

IV. PERFORMANCE

A. Metrics description

In order to assess the performances of the optimization, four metrics are used. In his study of the one dimensional problem, Ruckle identified the uniform distribution of paths along a line as the optimal strategy. This strategy corresponds to maximizing the entropy of the probability distribution. Based on this idea, the probability of crossing along the median between departure and arrival is computed. The *entropy* is used as the first metric, it is high for deceptive routes and close to zero for non-deceptive ones. The idea of the optimal strategy being uniformly distributed on the environment led to the creation of another metric, the *spreading* of the route, which measures the portion of the environment covered by the route as a fraction of the total surface of the environment. For any ambush area the probability of entering this area can be computed as the sum of inward flows. If the area is reached with probability $p > p_{min}$, the surface of this area is computed in the spreading. The *strategic outcome* V as computed in (1) is one of the most important metrics. Recall that V accounts for the expected losses for BLUE during one trip. It reflects how well the route can be optimized to reduce these losses. The *energy* metric is defined as $E = \sum_{i,j|(i,j) \in \mathcal{E}} p_{ij} \|e_{ij}\|$. This metric represents the divergence of the stochastic strategy from the shortest path between origin and destination.

B. Results

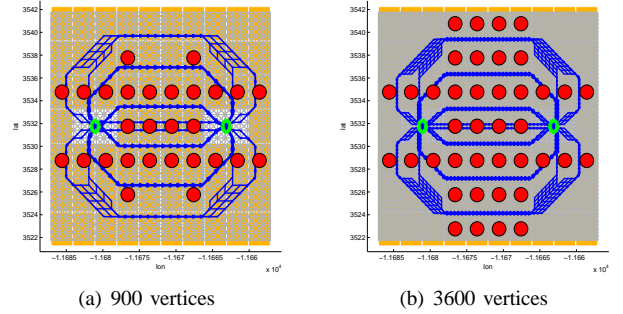


Fig. 3. Comparison of the results obtained for different network density with the second network construction method. The departure-arrival median is divided in seven areas, which leads to seven different paths. The routes in (a) and (b) are sensibly alike, illustrating the convergence of our solution when the size of the network increases.

Figure 3 illustrates one of the main results of this paper. For each method, BLUE's strategy is converging to a distribution that depends solely on RED's reach and the local outcome. The reach of RED in the following environment is such that the median between the departure and arrival points is divided into seven sections. This leads to an optimal theoretical outcome of $\frac{1}{7} \simeq 0.1456$ and hence a theoretical optimal entropy of $\sum_{i=1}^7 -\frac{1}{7} \log(\frac{1}{7}) = \log(7) \simeq 1.94$.

The results of the sensitivity analysis are represented in Figure 4. The metrics defined above are computed for each network construction method, using the simplex or the interior point optimization algorithms, with or without energy optimization, are compared. The energy optimization was introduced to avoid the presence of cycles in the solutions. However, the parameter λ needs to remain small compared to 1, otherwise it would interfere with the strategic optimization goal. Figure 4 display the metrics as function of the number of vertices in the network representation of the environment. Comparing them for the different network construction methods and optimization algorithms support the choices to make as to which one are more suited to build an optimal strategy. For the random sampling method, the metrics displayed correspond to the average over ten runs for every network size N .

In Figures 4(a),4(b), the strategic outcome converges for each optimization method and construction choice of the network. Regular networks, such as `uni8` and `uniD` converge faster with or without energy optimization. As soon as there exists a path for each of the seven distinct sections in the environment described above, the solutions are close to optimality. A similar behavior is observed in Figures 4(g),4(h). In Figures 4(c),4(d), two behaviors are observed depending on whether the energy is optimized. On average, more spreading is present when constructing a random network. When optimizing the energy, the regular networks exhibit less spreading, and the spreading converges to a smaller value. Also, the simplex method tends to provide less spread solutions when there is no energy optimization. In Figures 4(e),4(f), the random networks are suboptimal in terms of energy repartition. However the energy optimization factor forces the convergence. For regular networks, the energy factor converges quickly. For $\lambda = 0$, the space of solutions is larger than for $\lambda = 10^{-4}$. Therefore, it is in the first case that the simplex and interior point methods give the most different results. The spreading and energy metrics tend to favor opposite solutions.

Figure 4(e) shows that the lack of such an optimization can lead to an expected length of travel that grows as $O(N^2)$. This property justifies the need for a non-zero energy coefficient λ . Comparing the network construction methods, `uni8` and `uniD` offer equivalent results while random sampling is more energy consuming according to Figure 4(f) results in lesser entropies on the distribution (Figures 4(g),4(h)). Note that the max theoretical entropy of $\log(7) \simeq 1.94$ is reached by `uni8` and `uniD`. The metrics and methods used converge after a given density of nodes in the network, for a fixed size of the environment. This corresponds to the presence of an edge between each pair of nodes where ambush areas are located. The limit observed in Figures 4(a),4(b) is reached when the network density becomes large enough so that BLUE's distribution can use a critical number of ambush areas, which is seven here. Hence there exists a sufficient network density for the representation of a two dimensional continuous environment. Interestingly, beyond critical network size, the solutions become similar to the *Hunter and Bird*

Game described in [13]. There seems to be an analogy between the number of significant paths in the optimal route and the set of homotopy classes that dividing the median would result in.

Overall better results are obtained using the Interior Point algorithm on a uniformly sampled network connected through Delaunay triangulation and with an energy coefficient $\lambda = 10^{-4}$. This technique will be applied to unstructured environments in the remainder of this paper.

V. APPLICATIONS

A. General framework description

The environment is analyzed for three different purposes. The first objective consists in completing the existing road network with an off-road network to cover both structured and unstructured parts of the environment. The second is to identify relevant geographical areas for ambushes. The third is to compute the local outcome map. A correlation is established at each location between the possible outcome of an ambush and the strategic characteristics of the local environment around this ambush.

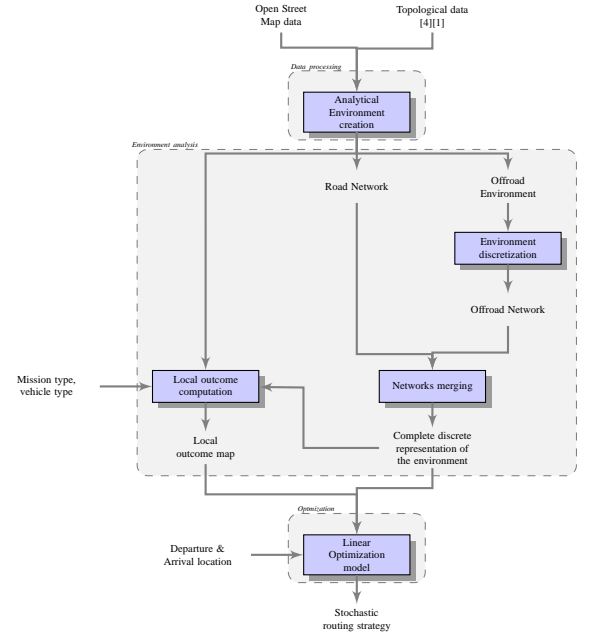


Fig. 5. Framework information pipeline.

To achieve these objectives, data of various natures is used. The elevation information about the environment is collected through data from the Shuttle Radar Topography Mission [4], which are publicly available [1]. SRTM data is distributed in two precision levels: SRTM1 (for the U.S. and its territories and possessions) with data sampled at one arc-second intervals in latitude and longitude, and SRTM3 (for the entire world) sampled at three arc-seconds. Three arc-second data are generated by three by three averaging of the one arc-second samples. An assessment of the SRTM data precision is given in [11]. While it is not precise enough for real-time high precision

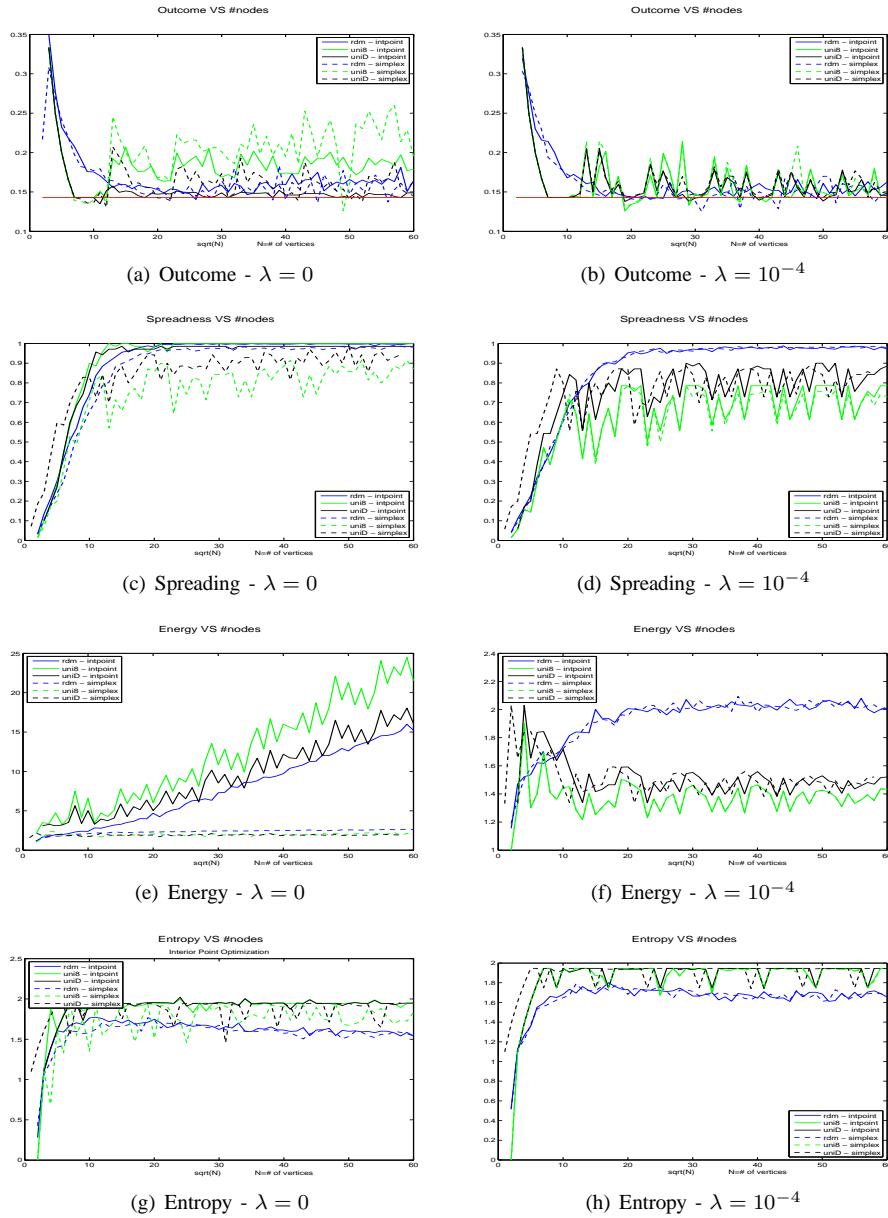


Fig. 4. Metrics convergence. The metrics are plotted as function of \sqrt{N} . All metrics convergence (except the energy metric) for the three methods for both values of λ . rdm creates slightly less optimal strategies regarding the entropy for which uni8 and uniD converge to the theoretical value of $\log 7$. rdm also has a slower convergence rate regarding the outcome.

path planning, which is not our present concern, it gives sufficient information for route planning purposes on scales of approximately one to tens of kilometers. A typical scenario takes place over 25-30 km in width, with points sampled every hundred meters. The second type of data used in this framework is Open Street Map[5] data. Open Street Map is a global, collaborative effort to create an open map of the world. The data aggregates semantic and usage information about the environment infrastructure such as the type of roads (interstate, primary, secondary, etc), authorized direction (one-way or two-ways), maximum speed, street names or information about the buildings.

The global framework pipeline is presented in Figure 5. The

different sources of data help create the analytic environment. This feeds into three layers: on the one hand the local outcome map, on the other hand the off-road environment and road information, which merge into the network to be used for the optimization. These three combined lead to the optimization model resulting into the stochastic optimal route.

As opposed to classical approaches, which require a number of manual operations and fine tuning, the methodology proposed offers an easy straightforward way to find routing strategies. Given the coordinates of the departure and arrival nodes and the type of vehicle, a local outcome map is generated and the optimal stochastic route is obtained.

B. Risk factors

In order to extend the purely theoretical game theory problem and make it more realistic, interviews of former Army personnel were performed. The focus of these interviews was to identify the factors with the highest influence on the local outcome that could be incorporated into the optimization problem formulation. Significant attention was devoted to the identification of the different factors influencing the local outcome.

The relevant information can be categorized into three different areas represented in Figure 6. First, the information related to the environment, whether described by topological or geospatial data, is decisive. This type of information covers data on elevation, roadways, waterways as well as vegetation and terrain type and quality. Second, the strategical information, such as RED's and BLUE's resources (position, amount, type) as well as previous ambushes, needs to be taken into account. Third, the human factors correspond to other intelligent agents present in the geographic region of interest that do not participate in the game but whose behavior may provide additional information. The traffic, the general behavior of a crowd or the amount of people in a given location are examples of such information. Note that these categories are correlated. For example, for a given type of vehicle, the maximum speed depends on the topology of the environment. A car might be faster than a human on a road but not on steep terrain. Such dependencies are accounted for in the information processing phase.

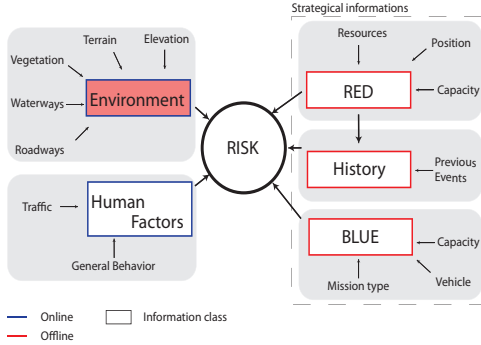


Fig. 6. Different factors to consider for local outcome computation at a given location.

One of the key factors that influence the local outcome of an ambush at a given location is the ability to apprehend a situation as early as possible. Such ability encompasses detecting changes, identifying unusual behavior or setting. This may rely on vision capabilities, in particular sight range and quality. Another key factor consists in having means to react quickly and possibly leave the location. Hence, the player's local maximum speed has a direct impact on the local outcome of an ambush at a given location.

The local outcome map heavily depends on the mission strategical information. The same topological environment, with the same resource allocation between the players, pro-

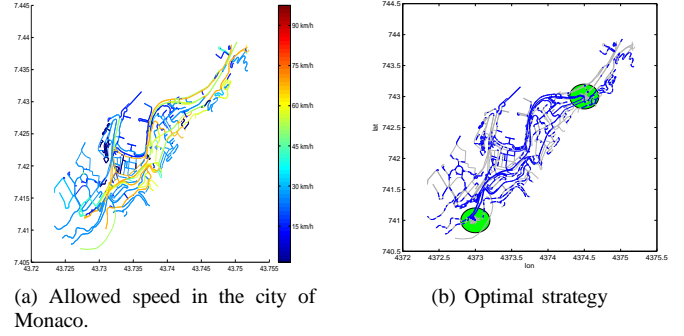


Fig. 7. Example of solution for the road network of the city of Monaco imported from OSM data.

duces a different outcome map depending on the type of mission. Three typical missions are briefly described below as examples. A scouting mission means that BLUE wants to travel from its departure point to its destination point, without being seen at all. The local outcome is then binary : if BLUE is seen, it loses. A transport mission aims at moving goods or assets from A to B. The local outcome is then directly linked to the proportion of the convoy that arrives safely to B. The opposite mission in terms of local outcome is a patrolling mission. The goal is to travel from A to B but a potential encounter would result in a win for BLUE. The local outcome is therefore positive if such an encounter occurs, which is not the case in a transport mission.

C. Examples

In the previous section, several factors influencing the local outcome were identified. The remainder of the paper focuses on examples of applications. The local outcome α is now defined as a parameter depending only on the maximum speed, and inversely proportional to it. A slower speed implies a higher local outcome, and reciprocally. The environmental data is used to find the maximum speed at a given position.

1) Road network:

For the example of the city of Monaco, the portion of unstructured environment is negligible. Therefore the model of ambushes at nodes is used. On road networks the maximum speed is set to be the maximum allowed speed, or the maximum safe driving speed. The local outcome at a given node is proportional to the mean time spent on the edges leading to this node. The speed profile of Monaco is presented in Figure 7(a). The corresponding optimal solution is displayed in Figure 7(b). The departure and arrival nodes are located on opposite sides of the city.

As one might expect, the optimal strategy returned by the framework uses with much higher probability the portions of the network where the speed limit is higher. However, the entire network is used by the routing strategy.

2) Unstructured Environment:

On the unstructured part of the environment, the maximum speed is a combined function of the different topological factors and information about BLUE's resources such as its means of locomotion. A three dimensional representation of




Fig. 8. 3D representation of the environment near Fort Irwin, CA. Imported from STRM data.

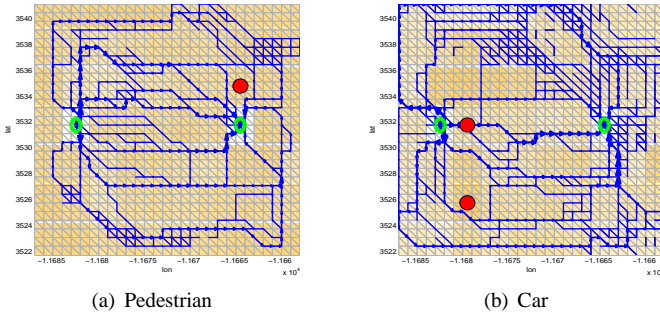


Fig. 9. Optimal routing strategy for a pedestrian and a car near Fort Irwin, CA. The pedestrian is slow and therefore its local outcome is high everywhere. The car's outcome is low in the lowlands (clear color) because it can go much faster. The pedestrian is comparatively less penalized in the hills than the car because its speed is less dependent on the terrain topology.

the environment is displayed in Figure 8, where lowlands and hills are present.

The resulting optimal strategies for either a pedestrian or a car are displayed in Figures 9(a) and 9(b). The variations of topology in the environment result in local outcome maps that are different for two different means of transportation. The pedestrian being much slower than the car in the lowlands, the local outcome is larger. On the contrary, it is much less penalized by the hills than the car is. Therefore it is more advantageous for the pedestrian to travel through the hills than it is for the car.

The idea that some resources for BLUE would be more adequate given a specific environment can lead to an interesting expansion to multimodal routing strategies. Given a set of transportation options, a player could change part of its dynamics during its travel in order to reduce the optimal outcome.

Overall the framework developed is efficient. Since it requires few inputs from the user, it can be easily used to compare different routing strategies on various environments. The large amount of SRTM and OSM data available makes it applicable to almost anywhere in the world.

VI. FUTURE WORK

From the theoretical standpoint, several aspects could be further investigated. The convergence observed for the routes among all methods suggests the existence of a limit distribution. The authors would like to propose a continuous model of

the game at hand instead of the dense discretization developed in this paper. This would allow to identify the limit distribution at once and maybe draw a parallel between the distinct paths among routes and homotopy classes.

A number of applications have been envisioned. The only factor influencing the local outcome is currently the maximum speed at a given location. The fidelity of the model could be improved by taking into account more risk factors in more detail. The purpose of this framework is to provide a hands-on tool for route prediction. However the merging of structured and unstructured environments still requires manual tuning. This merging process will be automated more thoroughly in the upcoming work. Finally, it might be interesting to increase the precision of the environment description by one or two orders of magnitude. Doing so requires determining how to accurately forecast the maximum safe speed at a given location. The control of the vehicle hence becomes part of the optimization problem.

VII. CONCLUSION

The work developed build on a previous research effort in which the environment was structured and no study of the local outcome of an ambush inference was performed. The present research defines the opponents reach as an important parameter of the environment representation. The local outcome becomes more meaningful because of its geographical significance. The possible losses for the system at a given area in case of ambush are computed through the identification of the different factors influencing this outcome. These factors are related to the environment and to the agents' resources. Comparing different network construction methods and different linear optimization algorithms, efficient techniques were identified to elaborate ambush avoidance strategies on unstructured environments. Results highlight the existence of a sufficient network density to represent these environments. The correlation between the reach of the ambushing player and the optimal outcome of the game is illustrated and a "representative distribution" emerges. Finally a comprehensive framework is elaborated. It takes as input topological data and, given a pair of points and a vehicle type, provides an optimal stochastic routing strategy between these two points. Examples of application are tested for structured and unstructured environments.

ACKNOWLEDGEMENT

This work was supported by the Army Research Office under MURI Award W911NF-11-1-0046.

REFERENCES

- [1] URL dds.cr.usgs.gov.
- [2] *Computing the Optimal Mixed Strategy for Various Ambush Games*, July 2005. doi: 10.1177/0278364910369189.
- [3] Emmanuel Boidot and Eric Feron. Planning random path distributions for ambush games in unstructured environments. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2012.

- [4] Tom G. Farr, Paul A. Rosen, Edward Caro, Robert Crippen, Riley Duren, Scott Hensley, Michael Kobrick, Mimi Paller, Ernesto Rodriguez, Ladislav Roth, David Seal, Scott Shaffer, Joanne Shimada, Jeffrey Umland, Marian Werner, Michael Oskin, Douglas Burbank, and Douglas Alsdorf. The shuttle radar topography mission. *Reviews of Geophysics*, 45(2):n/a–n/a, 2007. ISSN 1944-9208. doi: 10.1029/2005RG000183.
- [5] Mordechai (Muki) Haklay and Patrick Weber. Open-streetmap: User-generated street maps. *IEEE Pervasive Computing*, 7(4):12–18, October 2008. ISSN 1536-1268. doi: 10.1109/MPRV.2008.80.
- [6] R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Dover Books on Mathematics. Dover Publications, 2012. ISBN 9780486135984.
- [7] Michal Jakob, Ondrej Vanek, and Michal Pechoucek. Using agents to improve international maritime transport security. *IEEE Intelligent Systems*, 26(1):90–96, january 2011. ISSN 1541-1672.
- [8] Farmey A. Joseph. Path-planning strategies for ambush avoidance. Master’s thesis, MIT, 2005.
- [9] Sertac Karaman and Emilio Frazzoli. Incremental sampling-based algorithms for a class of pursuit-evasion games. In *WAFR*, volume 68 of *Springer Tracts in Advanced Robotics*, pages 71–87. Springer, 2010. ISBN 978-3-642-17451-3.
- [10] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944. ISBN 0691119937.
- [11] E. Rodriguez, C. S. Morris, J. E. Belz, E. C. Chapin, J. M. Martin, W. Daffer, and S. Hensley. *An assessment of the SRTM topographic products*. Jet Propulsion Laboratory, Pasadena, California, 2005.
- [12] William Ruckle, Robert Fennell, Paul T Holmes, and Charles Fennemore. Ambushing random walks i: Finite models. *Operations Research*, 24(2):314–324, 1976.
- [13] William H. Ruckle. Geometric games of search and ambush. *Mathematics Magazine*, 52(4):pp. 195–206, 1979. ISSN 0025570X.
- [14] Matteo Salani, Gaëtan Duyckaerts, and Peter Goodings Swartz. Ambush avoidance in vehicle routing for valuable delivery. *Journal of Transportation Security*, 3(1): 41–55, 2010.
- [15] Ondrej Vanek, Michal Jakob, Ondrej Hrstka, and Michal Pechoucek. Agent-based model of maritime traffic in piracy-affected waters. *Transportation Research Part C: Emerging Technologies (to appear)*, 2013. 10.1016/j.trc.2013.08.009.